ELSEVIER

Neurocomputing 26–27 (1999) 901–910

NEUROCOMPUTING

www.elsevier.com/locate/neucom

Attention, intention and salience in the posterior parietal cortex

Kathleen Taylor*, John Stein

University Laboratory of Physiology, University of Oxford, Parks Road, Oxford, OX1 3PT UK

Accepted 18 December 1998

Abstract

Selecting visual targets for saccadic eye movements is a vital step in sensorimotor processing. This selection is made on the basis of target salience: a saccade tends to be made to the most interesting part of the visual field. Both bottom-up and top-down processes have been postulated to contribute to salience, but the exact mechanisms by which some areas of the visual field are rendered to be more interesting than others remain unclear. We propose a computational model to address this issue and investigate its significant features, relating it to neuroanatomy and to aspects of cognitive function.

Keywords: Attention; Salience; Parietal; Lateral inhibition

1. Introduction

In any visual scene there are large number of potential targets for saccadic eye movements. In practice, these are weighted according to their salience – how interesting and relevant they are to the organism – with only the most salient being selected as actual targets for saccades. Therefore there must be some representation of salience in the brain. It is known that the areas of cortex which are involved in the generation of saccades encode representations of spatial location. This suggests that some combined representation – encoding spatial locations in terms of their salience – may be utilised. Such a salience representation (SR) in the brain could represent spatial locations in terms of how salient those locations themselves are: on the basis of learned cues for

* Corresponding author. E-mail: kathleen.taylor@physiol.ox.ac.uk.

0925-2312/99/$ – see front matter
PII: S0925-2312(98)00135-0

example. It could also weight the locations according to the salience of objects currently or previously situated at (or anticipated to be at) those locations. The SR could thus play a major role in the generation of eye movements, from fast single saccades to antisaccades or memory-guided movements.

The factors contributing to salience are numerous and complex. They certainly include "bottom-up" visual parameters, such as luminance contrast, size and colour [20]. However, the fact that saccades can be directed to areas of visual space on the basis of prior expectation, memory or emotional association indicates that "top-down" influences must also contribute to what makes a target salient [6,14,17]. This raises an important question: how can all these possible sources of information about the visual world be combined to select targets for eye movements?

1.1. Generation of the SR

It has been suggested that bottom-up inputs are "filtered" or "structured" by top-down inputs [19]. The latter provide an up-to-date high-level "hypothesis" about what is expected, with which the former are continually compared. As it stands this hypothesis implies that both are required at any given moment. However this is clearly not the case. Novel objects, which by definition evoke no memories or expectations, and which may have no emotional associations, nevertheless, evoke not just orienting responses but also voluntary exploratory saccades. Conversely, expectation or memory can invest an otherwise unremarkable part of the visual scene with significance for the organism. Thus the SR must be able to function over a wide variety of possible inputs, each of which may or may not be present at any given time. Furthermore, since auditory and mnemonic cues can be used to prime particular locations, the inputs to an SR must also be multimodal.

So far only the inputs to the SR have been discussed. However, if the SR is to play a part in the selection of targets for eye movements, further organisation of those inputs is required. Crudely, any limb can only make one movement at a time. Multiple areas activated by potential targets send convergent information to motor control areas; that is, the potential targets can be said to compete for control of motor outputs, with the most salient target winning. Because of this, and because of the temporal limits of short-term memory for motor plans, which means that only a few movements to targets can be stored, it is necessary to reduce the number of potential targets. Furthermore, this reduction should enhance the selected minority relative to the rest of the SR, since this will facilitate their ability to compete for access to output cortical and subcortical control structures.

1.2. The posterior parietal cortex as a candidate SR

It has been proposed that such a salience representation is generated by the posterior parietal cortex [22]. In particular, cells in the lateral intraparietal sulcus (area LIP) appear to encode spatial locations [5,13]. More recently, Gottlieb et al. [16] found evidence of a representation of only the most salient visual targets in this area. LIP is thought to be important in saccadic targeting [8]. Many PPC cells also

show multimodal responses [2,18]. Computationally, the simplest encoding of saliency is by neural activation: more active areas represent more salient spatial locations. This is the encoding adopted in this study.

The PPC in general, and LIP in particular, is ideally placed to serve as a multimodal co-ordinating centre. It is interconnected with a wide range of cortical and subcortical areas [21], including visual [1,9], auditory [12], limbic [7] and frontal areas [11]. These converging inputs suggest that the PPC receives information not only about many different aspects of the sensory environment, but also about intended movements [3] as well as the drive state of the organism. This makes the PPC a strong candidate for the SR.

Two other roles proposed for the PPC are in motor intention [4] and visuospatial attention [15]. Any SR description of PPC function must therefore give an account of the relationship between attention, intention and salience.

Attention is typically described as serial, voluntary and fast (tens to hundreds of milliseconds). It is multimodal, sensorimotor, and involves target selection; as does the as yet unspecified mechanism which generates the SR. Occam's razor suggests that these two mechanisms may be the same. Attention can therefore be thought of as a function operating to generate the SR.

On the other hand, Andersen has proposed that the PPC encodes the intention to move. What does this mean? Intention is dissociable from movement, just as attending to an object dissociates from saccading to it. Intention is also voluntary, like attention. Again a merger seems reasonable. Attention in its sensory aspect mediates the selection of visual targets. Attention in its motor aspect mediates the selection of motor goals, and this is what we call intention.

Evidence suggests that attention may involve a process of spatially localised relative enhancement [10] by which potential targets compete. The aim of our study was to build a computational model of the SR which displays such enhancement. Some kind of lateral inhibition is the most likely mechanism. However, the exact nature of the inhibition function is unclear. Proposals have included: that the inhibition falls off as a Gaussian, a difference of Gaussians, a step, or that it is uniform. We therefore compared these four functions to see which, if any, evoked enhancement.

2. Methods

Simulations were written in MATLAB versions 4.2/5.0/5.1, using a grid of 10×10 units with fixed inhibitory connections to represent the SR. Salience was encoded by a scalar intensity value, the activation A, given to each unit in the grid. (In practice, the value of the salience scalar will be set by the inputs to the PPC, which include "sensory" (e.g. visual, auditory), "motor" (corollary discharge) and "cognitive" (e.g. motivational and mnemonic information from limbic and prefrontal areas). An input pattern S was applied at the start of each trial, simulating the sum of inputs to the units. This pattern consisted of a Gaussian "blob" (peak value 1, standard deviation 1, centred on a pseudo-randomly selected location on the grid) with added uniform random noise in the range [0, 0.6]. At each time step t each unit's activation A_j (units

were indexed $j = 1, \ldots, N$, where N is the number of units in the grid) was synchronously adjusted according to a discrete update equation, up to $t = 25$. For each unit

$$A_{(0)} = S; \qquad A_{(t+1)} = \max((A_{(t)} - I_{(t)}), 0),$$

where $I_{(t)}$ is the total inhibition acting on the unit at time t. Units were bounded in the range $[0, +\infty]$. Except for the uniform case, inhibition between units depends on their relative locations within the grid. The nature of this dependence is governed by a "width" parameter W. We investigated four lateral inhibition functions (uniform, step, Gaussian and difference of Gaussians) to see which was most effective at producing localised enhancement. For each function F, the resultant inhibition of a unit indexed i at a given time step was given by

$$I_i = \sum_j (F_{ij} A_j),$$

where F_{ij} is the value of the inhibition function between units j and i. All inhibition was symmetrical. W is the width parameter. The four inhibition functions are

Uniform:

$$F_{ij} = \frac{1}{N}.$$

Step:

$$F_{ij} = \frac{1}{N}, \quad \Delta_{ij} \preceq W, \qquad F_{ij} = 0, \quad \Delta_{ij} \succ W.$$

Gaussian:

$$F_{ij} = \frac{1}{\sqrt{2\pi W^2}} \exp\left(- \frac{(x_i - x_j)^2 + (y_i - y_j)^2}{2W^2} \right).$$

Difference of Gaussians (DOG):

$$F_{ij} = \frac{1}{\sqrt{2\pi W_1^2}} \exp\left(- \frac{(x_i - x_j)^2 + (y_i - y_j)^2}{2W_1^2} \right) - \frac{1}{\sqrt{2\pi W_2^2}} \exp\left(- \frac{(x_i - x_j)^2 + (y_i - y_j)^2}{2W_2^2} \right).$$

For a DOG centred on the receiving unit i, whose coordinates in the 2-D layer are $[x_i, \ y_i]$, the widths for the two Gaussians are given by W_1 and W_2. We used $W_1 = 5W_2$.

To compare the results quantitatively, we developed a measure of image "complexity", the performance measure Φ, to assess enhancement. Φ reflects how the difference between the activation of the "best" unit, a_c, and that of other units

changes over time. It is 0 if all the units have the same activation, and is 1 if all the units except one have the same value:

$$\Phi = \left(\sum_{j}^{N} (a_c - a_j)/a_c(N-1) \right), N > 1, \quad a_c > 0.$$

For each inhibition function, except for the uniform function which has no width parameter W, trials were run for a range of widths $W = 1$–25. Enhancement was assessed by visualisation and by calculating Φ.

3. Results

Fig. 1 plots activations of the units for a trial in which enhancement occurred. F is the step function, at a sample width $W = 10$. Each plot shows unit activations at a given time t, coded by activation value (white squares represent the highest and black squares the lowest activations). Fig. 2 shows activations for a trial where no enhancement occurred. F is the DOG function, $W = 10$. The step function achieves enhancement within a few time steps; the DOG function does not.

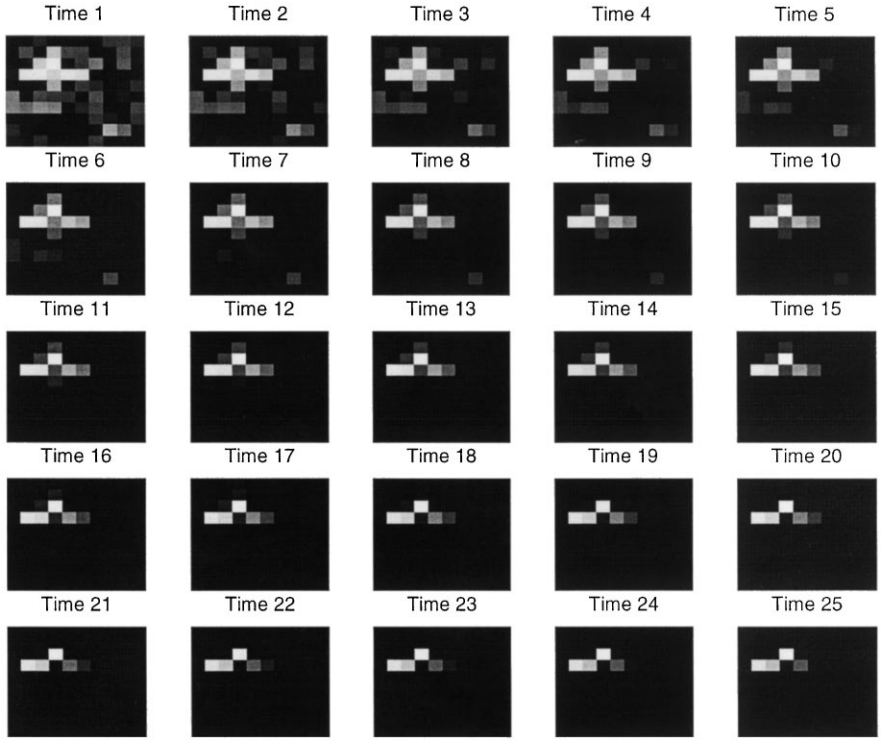

Fig. 1. Enhancement in a 10×10 activation matrix using the step inhibition function.

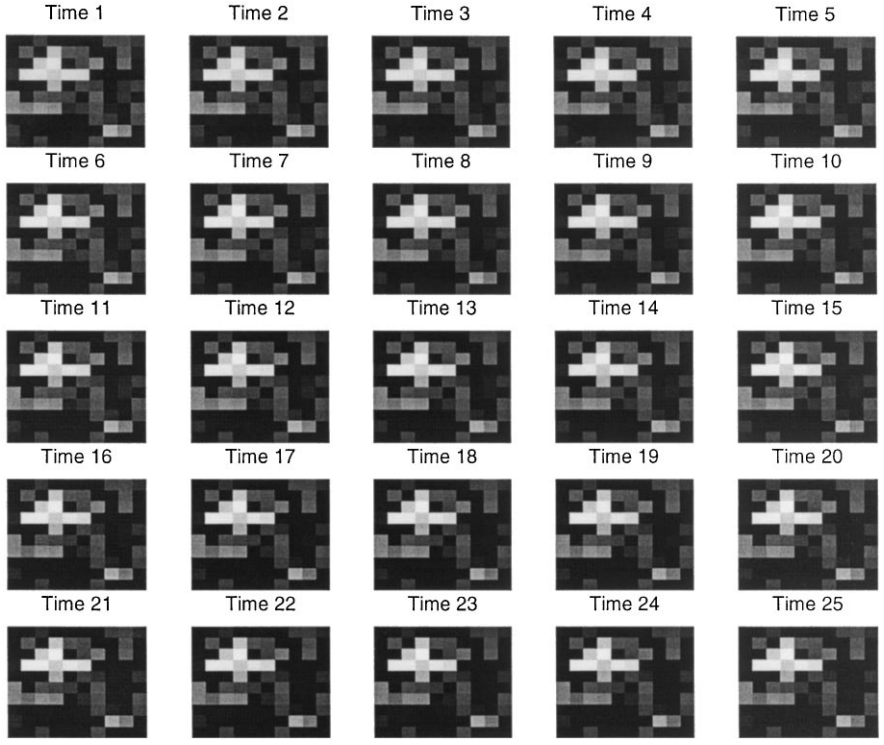

Fig. 2. Enhancement in a 10×10 activation matrix using the DOG inhibition function.

Fig. 3 plots Φ against time at a selection of widths $W = 1, 2, 5, 10, 25$. Each sub-plot shows the results for a different inhibition function F. The step function performed best. The optimum width appears to be $W = 10$ units; beyond this the performance reached a plateau. Surprisingly, the other inhibition functions failed to achieve enhancement.

These results indicate that enhancement is only achieved by the step function. The uniform function is over-simplistic. However, the Gaussian and DOG, which might have been expected to perform well, appear in fact to perform less effectively than the simpler step function.

We also ran two further simulations. The first investigated the effect of varying matrix connectivity for the best-performing step function, $W = 10$. This was done by setting each inter-unit connection to "ON" or "OFF" with probability C in the range 0 (no connections) to 1 (full connectivity). At each value of C (expressed in percentage terms in Fig. 4, such that 100% = full connectivity), we measured the number of time steps taken by the simulation to reach a performance level of $\Phi = 0.9$ (90% criterion) and $\Phi = 0.95$ (95% criterion). Fig. 4 shows the results. The solid horizontal line indicates the cut-off point; values above this level never reached criterion during the 25 time steps of the trial. In general, the lower the connectivity, the longer the

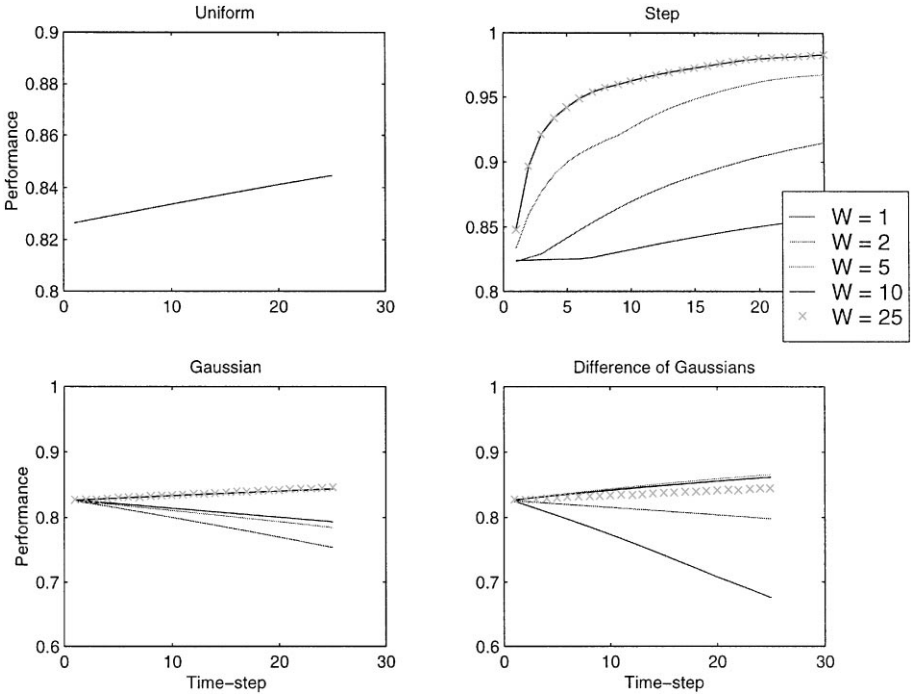

Fig. 3. Performance of the inhibition functions over time for a range of widths.

simulation takes to reach criterion. At connectivities of $C \geq 22\%$, the simulation takes fewer than 10 steps to reach the 90% criterion; at higher connectivities ($C \geq 63\%$) simulation takes fewer than 15 steps to reach the 95% criterion. If we assume that each time step in the simulation corresponds to a processing step in the cortex (e.g. a synaptic transfer, of the order of 5 ms), then the enhancement effect takes shape well within the timescale of ten to hundreds of ms required for attention.

The second simulation removed the bounding of unit activations at zero or higher. This abolished the enhancement effect for all inhibition functions at all widths, suggesting that zero-bounded activations are essential for relative enhancement to occur. Since the activations of units are taken to represent some monotonic function of neuronal firing rate (of a neuron or group of neurons), the activations cannot plausibly fall below zero.

4. Discussion

These simulations suggest that simple lateral inhibition is an effective means of achieving the enhancement required for target selection. Our algorithm achieves enhancement quickly, using a simple step inhibition function. We argue that such inhibition is biologically plausible and could be implemented by, for instance,

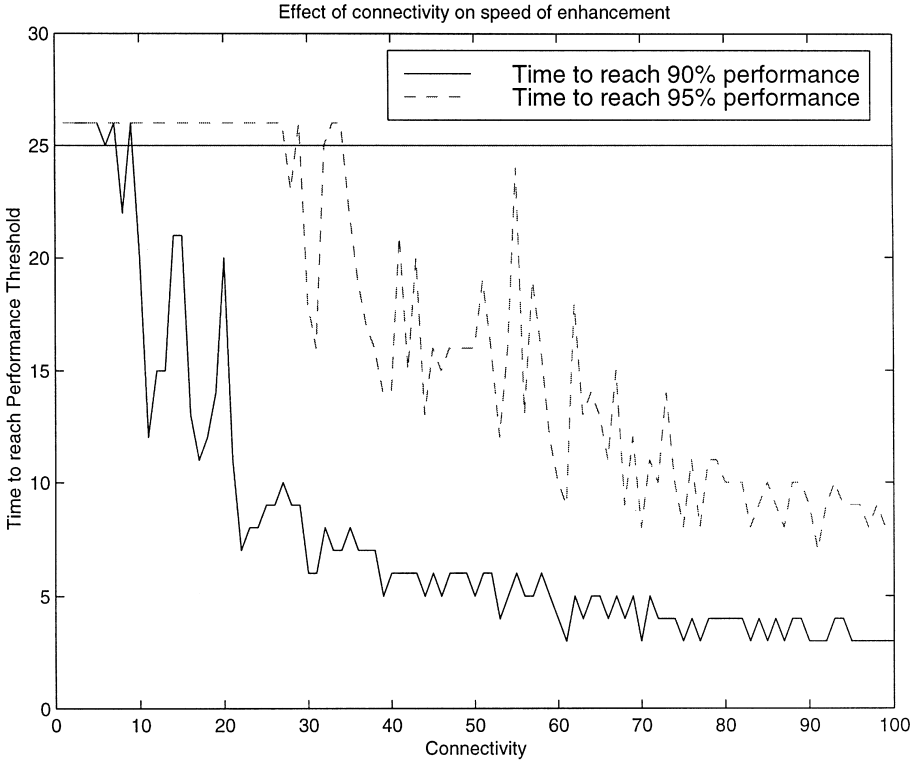

Fig. 4. Effect of connectivity on performance for the step inhibition function.

GABAergic interneurons connecting to a local cluster of surrounding cells. This simple mechanism, combined with the vast range of inputs which reach the posterior parietal cortex, could be sufficient to underlie the salience representation postulated to exist in the PPC. As we argued in the Introduction, this representation may be central to both attention and intention.

References

[1] R.A. Andersen, C. Asanuma, G. Essick, R.M. Siegel, Corticocortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule, J. Comp. Neurol. 296 (1990) 65–113.

[2] R.A. Andersen, R.M. Bracewell, S. Barash, J.W. Gnadt, L. Fogassi, Eye position effects on visual, memory, and saccade-related activity in areas LIP and 7a of macaque, J. Neurosci. 10 (1990) 1176–1196.

[3] R.A. Andersen, G.K. Essick, R.M. Siegel, Neurons of area 7 activated by both visual stimuli and oculomotor behavior, Exp. Brain Res. 67 (1987) 316–322.

[4] R.A. Andersen, L.H. Snyder, D.C. Bradley, J. Xing, Multimodal representation of space in the posterior parietal cortex and its use in planning movements, Ann. Rev. Neurosci. 20 (1997) 303–330.

[5] R.A. Andersen, D. Zipser, The role of the posterior parietal cortex in coordinate transformations for visual-motor integration, Can. J. Physiol. Pharmacol. 66 (1988) 488–501.
[6] W.F. Bacon, H.E. Egeth, Overriding stimulus-driven attentional capture, Percept. Psychophys. 55 (1994) 485–496.
[7] C. Baleydier, F. Mauguiere, Network organization of the connectivity between parietal area 7, posterior cingulate cortex and medial pulvinar nucleus: a double fluorescent tracer study in monkey, Exp. Brain Res. 66 (1987) 385–393.
[8] S. Barash, R.M. Bracewell, L. Fogassi, J.W. Gnadt, R.A. Andersen, Saccade-related activity in the lateral intraparietal area. I. Temporal properties; comparison with area 7a, J. Neurophysiol. 66 (1991) 1095–1108.
[9] D. Boussaoud, L.G. Ungerleider, R. Desimone, Pathways for motion analysis: cortical connections of the medial superior temporal and fundus of the superior temporal visual areas in the macaque, J. Comp. Neurol. 296 (1990) 462–495.
[10] M.C. Bushnell, M.E. Goldberg, D.L. Robinson, Behavioral enhancement of visual responses in monkey cerebral cortex. I. Modulation in posterior parietal cortex related to selective visual attention, J. Neurophysiol. 46 (1981) 755–772.
[11] C. Cavada, P.S. Goldman-Rakic, Posterior parietal cortex in rhesus monkey: I. Parcellation of areas based on distinctive limbic and sensory corticocortical connections, J. Comp. Neurol. 287 (1989) 393–421.
[12] J.F. Demonet, F. Chollet, S. Ramsay, D. Cardebat, J.L. Nespoulous, R. Wise, A. Rascol, R. Frackowiak, The anatomy of phonological and semantic processing in normal subjects, Brain 115 (1992) 1753–1768.
[13] P. Fattori, C. Galletti, P.P. Battaglini, Parietal neurons encoding visual space in a head-frame of reference, Boll. Soc. Ital. Biol. Sper. 68 (1992) 663–670.
[14] I.H. Fraser, G.L. Craig, D.M. Parker, Reaction time measures of feature saliency in schematic faces, Perception 19 (1990) 661–673.
[15] M.E. Goldberg, C.J. Bruce, Cerebral cortical activity associated with the orientation of visual attention in the rhesus monkey, Vision Res. 25 (1985) 471–481.
[16] J.P. Gottlieb, M. Kusunoki, M.E. Goldberg, The representation of visual salience in monkey parietal cortex, Nature 391 (1998) 481–484.
[17] N.D. Haig, High-resolution facial feature saliency mapping, Perception 15 (1986) 373–386.
[18] K.W. Koch, J.M. Fuster, Unit activity in monkey parietal cortex related to haptic perception and temporary memory, Exp. Brain Res. 76 (1989) 292–306.
[19] D. Mumford, On the computational architecture of the neocortex. II, The role of cortico-cortical loops, Biol. Cybernet. 66 (1992) 241–251.
[20] H.C. Nothdurft, The role of features in preattentive vision: comparison of orientation, motion and color cues, Vision Res. 33 (1993) 1937–1958.
[21] D.N. Pandya, B. Seltzer, Intrinsic connections and architectonics of posterior parietal cortex in the rhesus monkey, J. Comp. Neurol. 204 (1982) 196–210.
[22] J.F. Stein, The representation of egocentric space in the posterior parietal cortex, Controversies in Neuroscience Conf. I. Movement Control (1990, Portland, Oregon), Behav. Brain Sci. 15 (1992) 691–700.

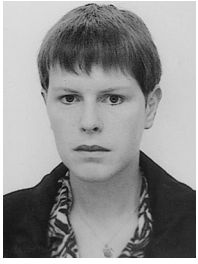

Kathleen Taylor studied physiology and philosophy at Oxford before doing a research M.Sc. in neuropharmacology at Stirling. She is currently writing up a D.Phil., at Oxford, in computational neuroscience.

John Stein read medicine at Oxford, then trained for clinical neurology at St. Thomas' Hospital in London, Oxford and Leicester. He was appointed to a Fellowship in Physiology at Magdalen College, Oxford, later followed by a Professorship. His research interests include parietal cortex, the visual guidance of movement, and dyslexia.